

[Home](#) | [Login](#) | [Logout](#) | [Access Information](#) | [Alerts](#) |

Welcome United States Patent and Trademark Office

SEARCH RESULTS**BROWSE****SEARCH****IEEE XPLORE GUIDE**

Results for "((very long instruction word) and simulat* and parallel and 'pipeline stage')<in>metadata..."
Your search matched 1 of 1589326 documents.



A maximum of 100 results are displayed, 25 to a page, sorted by **Relevance** in **Descending** order.

» Search Options[View Session History](#)**Modify Search**[New Search](#)((very long instruction word) and simulat* and parallel and 'pipeline stage')<in>metadata... **» Key****IEEE JNL** IEEE Journal or Magazine**IET JNL** IET Journal or Magazine**IEEE CNF** IEEE Conference Proceeding**IET CNF** IET Conference Proceeding**IEEE STD** IEEE Standard**Display Format:** Citation Citation & Abstract

1. **An instruction-level energy model for embedded VLIW architectures**
Sami, M.; Sciuto, D.; Silvano, C.; Zaccaria, V.;
Computer-Aided Design of Integrated Circuits and Systems, IEEE Transaction:
Volume 21, Issue 9, Sept. 2002 Page(s):998 - 1010
Digital Object Identifier 10.1109/TCAD.2002.801105

[AbstractPlus](#) | [References](#) | Full Text: [PDF\(368 KB\)](#) IEEE JNL
[Rights and Permissions](#)

[Help](#) [Contact Us](#) [Privacy & :](#)

© Copyright 2006 IEEE -

Indexed by
 Inspec®



[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)

Search: The ACM Digital Library The Guide

+ "very long instruction word" +simulat* +parallel +"pipeline s"

SEARCH

THE ACM DIGITAL LIBRARY

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used: **very long instruction word simulat parallel pipeline stage**

Found 44 of 204,472

Sort results by

[Save results to a Binder](#)

Try an [Advanced Search](#)

Display results

[Search Tips](#)

Try this search in [The ACM Guide](#)

Open results in a new window

Results 1 - 20 of 44

Result page: [1](#) [2](#) [3](#) [next](#)

Relevance scale

1 Effective compiler generation by architecture description

Stefan Farfeleider, Andreas Krall, Edwin Steiner, Florian Brandner

June 2006 **ACM SIGPLAN Notices**, Proceedings of the 2006 ACM SIGPLAN/SIGBED conference on Language, compilers and tool support for embedded systems LCTES '06, Volume 41 Issue 7

Publisher: ACM Press

Full text available: [pdf\(128.18 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Embedded systems have an extremely short time to market and therefore require easily retargetable compilers. Architecture description languages (ADLs) provide a single concise architecture specification for the generation of hardware, instruction set simulators and compilers. In this article, we present an ADL for compiler generation. From a specification, we can derive an optimized tree pattern matching instruction selector, a register allocator and an instruction scheduler. Compared to a hand- ...

Keywords: architecture description language, code generation, compiler generation

2 Static Placement, Dynamic Issue (SPDI) Scheduling for EDGE Architectures

Ramadass Nagarajan, Sundeep K. Kushwaha, Doug Burger, Kathryn S. McKinley, Calvin Lin, Stephen W. Keckler

September 2004 **Proceedings of the 13th International Conference on Parallel Architectures and Compilation Techniques PACT '04**

Publisher: IEEE Computer Society

Full text available: [pdf\(182.53 KB\)](#) Additional Information: [full citation](#), [abstract](#), [citations](#)

Technology trends present new challenges for processor architectures and their instruction schedulers. Growing transistor density will increase the number of execution units on a single chip, and decreasing wire transmission speeds will cause long and variable on-chip latencies. These trends will severely limit the two dominant conventional architectures: dynamic issue superscalars, and static placement and issue VLIWs. We present a new execution model in which the hardware and static scheduler ...

3 Embedded systems: applications, solutions and techniques (EMBS): DSPxPlore:

design space exploration methodology for an embedded DSP core

Christian Panis, Ulrich Hirnschrott, Gunther Laure, Wolfgang Lazian, Jari Nurmi

March 2004 **Proceedings of the 2004 ACM symposium on Applied computing SAC '04**

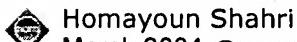
Publisher: ACM Press

Full text available:  pdf(324.12 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

High mask and production costs for the newest CMOS silicon technologies increase the pressure to develop hardware platforms useable for different applications or variants of the same application. To provide flexibility for these platforms the need on software programmable embedded processors is increasing. To close the gap concerning consumed silicon area and power dissipation between optimized hardware implementations and software based solutions, it is necessary to adapt the subsystem of the e ...

Keywords: DSPxPlore, design space exploration, embedded DSP

4 Queue Focus: DSP: On Mapping Algorithms to DSP Architectures 



Homayoun Shahri
March 2004 **Queue**, Volume 2 Issue 1

Publisher: ACM Press

Full text available:  pdf(644.31 KB)  html(34.78 KB) Additional Information: [full citation](#), [index terms](#)

5 Queue Focus: DSP: DSPs: Back to the Future 



W. Patrick Hays
March 2004 **Queue**, Volume 2 Issue 1

Publisher: ACM Press

Full text available:  pdf(1.80 MB)  html(34.87 KB) Additional Information: [full citation](#), [index terms](#)

6 LZW-Based Code Compression for VLIW Embedded Systems 

Chang Hong Lin, Yuan Xie, Wayne Wolf

February 2004 **Proceedings of the conference on Design, automation and test in Europe - Volume 3 DATE '04**

Publisher: IEEE Computer Society

Full text available:  pdf(398.84 KB) Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#), [review](#)

We propose a new variable-sized-block method for VLIW code compression. Code compression traditionally works on fixed-sized blocks and its efficiency is limited by the smallblock size. Branch blocks -- instructions between two consecutive possible branch targets -- provide larger blocks for code compression. We propose LZW-based algorithms to compress branch blocks. Our approach is fully adaptive and generates coding table on-the-fly during compression and decompression. When encountering a branc ...

7 Modeling and validation of pipeline specifications 



Prabhat Mishra, Nikil Dutt
February 2004 **ACM Transactions on Embedded Computing Systems (TECS)**, Volume 3
Issue 1

Publisher: ACM Press

Full text available:  pdf(198.92 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#), [review](#)

Verification is one of the most complex and expensive tasks in the current Systems-on-Chip design process. Many existing approaches employ a bottom-up approach to pipeline validation, where the functionality of an existing pipelined processor is, in essence, reverse-engineered from its RT-level implementation. Our validation technique is complementary to these bottom-up approaches. Our approach leverages the system

architect's knowledge about the behavior of the pipelined architecture, through a ...

Keywords: Modeling of processor pipeline, architecture description language, pipeline validation, pipelined processor specification

8 Embedded system architectures: Synthesizable HDL generation method for configurable VLIW processors

Yuki Kobayashi, Shinsuke Kobayashi, Koji Okuda, Keishi Sakanushi, Yoshinori Takeuchi, Masaharu Imai

January 2004 **Proceedings of the 2004 conference on Asia South Pacific design automation: electronic design and solution fair ASP-DAC '04 , Proceedings of the 2004 conference on Asia South Pacific design automation: electronic design and solution fair ASP-DAC '04**

Publisher: IEEE Press

Full text available: [pdf\(217.47 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#)

[Publisher Site](#)

This paper proposes a synthesizable HDL code generation method using a processor specification description. The proposed approach can change the number of slots and pipeline stages, and dispatching rule to assign operations to resources. In addition, designers can specify each instruction behavior using the specification language. A control logic, a decode logic, and a data path of VLIW processor are generated from the processor specification. Designers can explore ASIP design space using the pr ...

9 Emerging areas: A new look at exploiting data parallelism in embedded systems

Hillary C. Hunter, Jaime H. Moreno

October 2003 **Proceedings of the 2003 international conference on Compilers, architecture and synthesis for embedded systems CASES '03**

Publisher: ACM Press

Full text available: [pdf\(322.12 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper describes and evaluates three architectural methods for accomplishing data parallel computation in a programmable embedded system. Comparisons are made between the well-studied *Very Long Instruction Word (VLIW)* and *Single Instruction Multiple Packed Data (SIMpD)* paradigms; the less-common *Single Instruction Multiple Disjoint Data (SIMdD)* architecture is described and evaluated. A taxonomy is defined for data-level parallel archi ...

Keywords: DLP, DSP, ILP, SIMD, VLIW, architecture, data-level parallelism, embedded, media, processor, sub-word parallelism, telecommunications

10 Microprocessor architecture: Automatic generation of application specific processors

David Goodwin, Darin Petkov

October 2003 **Proceedings of the 2003 international conference on Compilers, architecture and synthesis for embedded systems CASES '03**

Publisher: ACM Press

Full text available: [pdf\(231.13 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

An application-specific instruction-set processor (ASIP) is ideally suited for embedded applications that have demanding performance, size, and power requirements that cannot be satisfied by a general purpose processor. ASIPs also have time-to-market and programmability advantages when compared to custom ASICs. The AutoTIE system simplifies the creation of ASIPs by automatically enhancing a base processor with application specific instruction set architecture (ISA) extensions, including instruct ...

Keywords: ASIPs, automatic instruction-set generation, configurable processors, extensible processors

11 HIBRID-SOC: a multi-core architecture for image and video applications

 S. Moch, M. Bereković, H. J. Stolberg, L. Friebe, M. B. Kulaczewski, A. Dehnhardt, P. Pirsch
September 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 2003 workshop on MEmory performance: DEaling with Applications , systems and architecture MEDEA '03**, Volume 32 Issue 3

Publisher: ACM Press

Full text available:  pdf(245.38 KB) Additional Information: [full citation](#), [abstract](#), [references](#)

The HiBRID-SoC multi-core architecture targets a wide range of application fields with particularly high processing demands, including general signal processing applications, video de-encoding, image processing, or a combination of these tasks. For this purpose, the HiBRID-SoC integrates three fully programmable processor cores and various interfaces on a single chip, all tied to a 64-Bit AMBA AHB bus. Its memory subsystem is particularly adapted to the high bandwidth demands of the multi-core a ...

12 Banked multiported register files for high-frequency superscalar microprocessors

 Jessica H. Tseng, Krste Asanović
May 2003 **ACM SIGARCH Computer Architecture News , Proceedings of the 30th annual international symposium on Computer architecture ISCA '03**, Volume 31 Issue 2

Publisher: ACM Press

Full text available:  pdf(142.29 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

Multiported register files are a critical component of high-performance superscalar microprocessors. Conventional multiported structures can consume significant power and die area. We examine the designs of banked multiported register files that employ multiple interleaved banks of fewer ported register cells to reduce power and area. Banked register files designs have been shown to provide sufficient bandwidth for a superscalar machine, but previous designs had complex control structures that w ...

13 Articles: Blurring Lines Between Hardware and Software

 Homayoun Shahri
April 2003 **Queue**, Volume 1 Issue 2

Publisher: ACM Press

Full text available:  html(23.93 KB) Additional Information: [full citation](#), [index terms](#)

14 HiBRID-SoC: A Multi-Core System-on-Chip Architecture for Multimedia Signal Processing Applications

Hans-Joachim Stolberg, Mladen Berekovic, Lars Friebe, Soren Moch, Sebastian Flugel, Xun Mao, Mark B. Kulaczewski, Heiko Klusmann, Peter Pirsch
March 2003 **Proceedings of the conference on Design, Automation and Test in Europe: Designers' Forum - Volume 2 DATE '03**

Publisher: IEEE Computer Society

Full text available:  pdf(307.90 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)

 Publisher Site

The HiBRID-SoC multi-core system-on-chip targets a wide range of application fields with particularly high processing demands, including general signal processing applications, video and audio de-/encoding, and a combination of these tasks. For this purpose, the HiBRID-SoC integrates three fully programmable processors cores and various interfaces

onto a single chip, all tied to a 64-Bit AMBA AHB bus. The processor cores are individually optimized to the particular computational characteristics ...

15 Flexible and Formal Modeling of Microprocessors with Application to Retargetable Simulation

Wei Qin, Sharad Malik

March 2003 **Proceedings of the conference on Design, Automation and Test in Europe - Volume 1 DATE '03**

Publisher: IEEE Computer Society

Full text available: [pdf\(185.73 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [citations](#), [index terms](#)

[Publisher Site](#)

Given the growth in application-specific processors, there is a strong need for a retargetable modeling framework that is capable of accurately capturing complex processor behaviors and generating efficient simulators. We propose the operation state machine (OSM) computation model to serve as the foundation of such a modeling framework. The OSM model separates the processor into two interacting layers: the operation layer where operation semantics and timing are modeled, and the hardware layer w ...

16 Design space exploration for embedded systems: Energy estimation and optimization of embedded VLIW processors based on instruction clustering

A. Bona, M. Sami, D. Sciuto, V. Zaccaria, C. Silvano, R. Zafalon

June 2002 **Proceedings of the 39th conference on Design automation DAC '02**

Publisher: ACM Press

Full text available: [pdf\(308.00 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Aim of this paper is to propose a methodology for the definition of an instruction-level energy estimation framework for VLIW (Very Long Instruction Word) processors. The power modeling methodology is the key issue to define an effective energy-aware software optimisation strategy for state-of-the-art ILP (Instruction Level Parallelism) processors. The methodology is based on an energy model for VLIW processors that exploits instruction clustering to achieve an efficient and fine grained energy ...

Keywords: power estimation, vliw architectures

17 Parallel and distributed systems and networking: Performance evaluation for a compressed-VLIW processor

Sunghyun Jee, Kannappan Palaniappan

March 2002 **Proceedings of the 2002 ACM symposium on Applied computing SAC '02**

Publisher: ACM Press

Full text available: [pdf\(498.01 KB\)](#)

Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

This paper presents a new ILP processor architecture called *Compressed VLIW* (CVLIW). The CVLIW processor constructs a sequence of long instructions by removing nearly all NOPs (No Operations) and LNOPS (Long NOPs) from VLIW code. The CVLIW processor individually schedules each instruction within long instructions using functional unit and dynamic scheduler pairs. Every dynamic scheduler in the CVLIW processor individually checks for data dependencies and resource collisions while scheduli ...

Keywords: CVLIW processor, ILP, VLIW, individual instruction scheduling

18 Optimizing a 3D image reconstruction algorithm: investigating the interaction between

the high-level implementation, the compiler and the architecture

Tom Vander Aa, Lieven Eeckhout, Bart Goeman, Hans Vandierendonck, Tanja Van Achteren, Rudy Lauwereins, Koen De Bosschere

January 2002 **Australian Computer Science Communications , Proceedings of the seventh Asia-Pacific conference on Computer systems architecture CRPIT '02**, Volume 24 Issue 3

Publisher: Australian Computer Society, Inc., IEEE Computer Society Press

Full text available: pdf(1.04 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [index terms](#)

Digital signal processing and multimedia workloads will be a dominant workload for computer based systems in the near future. In this paper, we evaluate the performance of an important media application, namely a relatively new 3D image reconstruction algorithm, on two platforms: a DSP processor (Texas Instruments TMS320C6701) and a high-performance general-purpose microprocessor (Alpha 21164). Prior to evaluating the performance of both architectural paradigms---very long instruction word (VLIW ...

Keywords: 3D image reconstruction, VLIW, program optimization, superscalar

19 Novel ideas: A design space evaluation of grid processor architectures

Ramadass Nagarajan, Karthikeyan Sankaralingam, Doug Burger, Stephen W. Keckler

December 2001 **Proceedings of the 34th annual ACM/IEEE international symposium on Microarchitecture MICRO 34**

Publisher: IEEE Computer Society

Full text available: pdf(1.29 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)
[Publisher Site](#)

In this paper, we survey the design space of a new class of architectures called Grid Processor Architectures (GPAs). These architectures are designed to scale with technology, allowing faster clock rates than conventional architectures while providing superior instruction-level parallelism on traditional workloads and high performance across a range of application classes. A GPA consists of an array of ALUs, each with limited control, connected by a thin operand network. Programs are executed b ...

20 Power-and Energy-Aware Computing: Comparing power consumption of an SMT and a CMP DSP for mobile phone workloads

Stefanos Kaxiras, Girija Narlikar, Alan D. Berenbaum, Zhigang Hu

November 2001 **Proceedings of the 2001 international conference on Compilers, architecture, and synthesis for embedded systems CASES '01**

Publisher: ACM Press

Full text available: pdf(315.81 KB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

In the DSP world, many media workloads have to perform a specific amount of work in a specific period of time. This observation led us to examine Simultaneous Multithreading (SMT) and Chip Multiprocessing (CMP) for a VLIW DSP architecture (specifically the Star*Core SC140), in conjunction with Frequency/Voltage scaling to decrease dynamic power consumption in next-generation wireless handsets. We study the resulting performance and power characteristics of the two approaches using simulation, co ...

Results 1 - 20 of 44

Result page: [1](#) [2](#) [3](#) [next](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads: Adobe Acrobat QuickTime Windows Media Player Real Player


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
Search: The ACM Digital Library The Guide

"very long instruction word" +simulat* +parallel +"pipeline s"

[THE ACM DIGITAL LIBRARY](#)
[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used: **very long instruction word simulat parallel pipeline stage**

Found 44 of 204,472

Sort results by publication date Save results to a Binder
 Display results expanded form Search Tips
 Open results in a new window

Try an [Advanced Search](#)
 Try this search in [The ACM Guide](#)

Results 21 - 40 of 44

 Result page: [previous](#) [1](#) [2](#) [3](#) [next](#)

Relevance scale

21 [Exploiting data forwarding to reduce the power budget of VLIW embedded processors](#)

M. Sami, D. Sciuto, C. Silvano, V. Zaccaria, R. Zafalon

 March 2001 **Proceedings of the conference on Design, automation and test in Europe DATE '01**

Publisher: IEEE Press

 Full text available: [pdf\(237.35 KB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: VLIW embedded architectures, forwarding, low-power, pipeline processors

22 [Polygon rendering on a stream architecture](#)

John D. Owens, William J. Dally, Ujval J. Kapasi, Scott Rixner, Peter Mattson, Ben Mowery
 August 2000 **Proceedings of the ACM SIGGRAPH/EUROGRAPHICS workshop on Graphics hardware HWWS '00**

Publisher: ACM Press

 Full text available: [pdf\(161.65 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The use of a programmable stream architecture in polygon rendering provides a powerful mechanism to address the high performance needs of today's complex scenes as well as the need for flexibility and programmability in the polygon rendering pipeline. We describe how a polygon rendering pipeline maps into data streams and kernels that operate on streams, and how this mapping is used to implement the polygon rendering pipeline on Imagine, a programmable stream processor. We compare our resul ...

Keywords: OpenGL, SIMD, graphics hardware, kernels, media processors, polygon rendering, stream architecture, stream processing, streams

23 [Exploiting ILP in page-based intelligent memory](#)

 Mark Oskin, Justin Hensley, Diana Keen, Frederic T. Chong, Matthew Farrens, Aneet Chopra
 November 1999 **Proceedings of the 32nd annual ACM/IEEE international symposium on Microarchitecture MICRO 32**

Publisher: IEEE Computer Society

Full text available: [pdf\(1.35 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)
[Publisher Site](#)

This study compares the speed, area, and power of different implementations of Active Pages [OCS98], an intelligent memory system which helps bridge the growing gap between processor and memory performance by associating simple functions with each page of data. Previous investigations have shown up to 1000X speedups using a block of reconfigurable logic to implement these functions next to each sub-array on a DRAM chip. In this study, we show that instruction-level parallelism, n ...

24 Value speculation scheduling for high performance processors

 Chao-Ying Fu, Matthew D. Jennings, Sergei Y. Larin, Thomas M. Conte
 October 1998 **ACM SIGOPS Operating Systems Review , ACM SIGPLAN Notices , Proceedings of the eighth international conference on Architectural support for programming languages and operating systems ASPLOS-VIII**, Volume 32 , 33 Issue 5 , 11

Publisher: ACM Press

Full text available: [pdf\(1.16 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Recent research in value prediction shows a surprising amount of predictability for the values produced by register-writing instructions. Several hardware based value predictor designs have been proposed to exploit this predictability by eliminating flow dependencies for highly predictable values. This paper proposed a hardware and software based scheme for value speculation scheduling (VSS). Static VLIW scheduling techniques are used to speculate value dependent instructions by scheduling them ...

Keywords: VLIW instruction schedulings, instruction level parallelism, value prediction, value speculation

25 Trace-driven studies of VLIW video signal processors

 Zhao Wu, Wayne Wolf
 June 1998 **Proceedings of the tenth annual ACM symposium on Parallel algorithms and architectures SPAA '98**

Publisher: ACM Press

Full text available: [pdf\(1.48 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: MPEG, VLIW, VSP, media processor, parallel architecture, parallelism, trace-driven scheduling, video applications

26 Realization of a programmable parallel DSP for high performance image processing applications

 Jens Peter Wittenburg, Willm Hinrichs, Johannes Kneip, Martin Ohmacht, Mladen Bereković, Hanno Lieske, Helge Kloos, Peter Pirsch
 May 1998 **Proceedings of the 35th annual conference on Design automation DAC '98**

Publisher: ACM Press

Full text available: [pdf\(2.35 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Architecture and design of the HiPAR-DSP, a SIMD controlled signalprocessor with parallel data paths, VLIW and novel memory design. The processor architecture is derived from an analysis of the target algorithms and specified in VHDL on register transfer level. A team of more than 20 graduate students covered the whole designprocess, including the

synthesizable VHDL description, synthesis, routing and backannotation as the development of a complete software development environment. The 175mm², 0.5µm ...

- 27 Performance analysis of tree VLIW architecture for exploiting branch ILP in non-numerical code**

 Soo-Mook Moon, Kemal Ebcioğlu
July 1997 **Proceedings of the 11th international conference on Supercomputing ICS '97**
Publisher: ACM Press
Full text available:  pdf(1.39 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

Keywords: branch code motion, conditional execution, generalized multiway branching, speculative code motion, tree VLIW architecture

- 28 Performance comparison of ILP machines with cycle time evaluation**

 Tetsuya Hara, Hideki Ando, Chikako Nakanishi, Masao Nakaya
May 1996 **ACM SIGARCH Computer Architecture News , Proceedings of the 23rd annual international symposium on Computer architecture ISCA '96**, Volume 24 Issue 2
Publisher: ACM Press

Full text available:  pdf(1.48 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Many studies have investigated performance improvement through exploiting instruction-level parallelism (ILP) with a particular architecture. Unfortunately, these studies indicate performance improvement using the number of cycles that are required to execute a program, but do not quantitatively estimate the penalty imposed on the cycle time from the architecture. Since the performance of a microprocessor must be measured by its execution time, a cycle time evaluation is required as well as a cy ...

- 29 The M-Machine multicomputer**

Marco Fillo, Stephen W. Keckler, William J. Dally, Nicholas P. Carter, Andrew Chang, Yevgeny Gurevich, Whay S. Lee
December 1995 **Proceedings of the 28th annual international symposium on Microarchitecture MICRO 28**
Publisher: IEEE Computer Society Press

Full text available:  pdf(1.29 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)

- 30 Speculative disambiguation: a compilation technique for dynamic memory disambiguation**

 A. S. Huang, G. Slavenburg, J. P. Shen
April 1994 **ACM SIGARCH Computer Architecture News , Proceedings of the 21ST annual international symposium on Computer architecture ISCA '94**, Volume 22 Issue 2
Publisher: IEEE Computer Society Press, ACM Press

Full text available:  pdf(1.09 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Ambiguous memory references have always been one of the main sources of performance bottlenecks. Many papers have addressed this problem using static disambiguation. These methods work extremely well when the memory access pattern is linear and predictable. However they are ineffective when the memory access pattern is nonlinear or when the access pattern cannot be determined statically. For these difficult problems, this paper

presents speculative disambiguation, a compilation technique for arc ...

31 Software support for speculative loads

 Anne Rogers, Kai Li

September 1992 **ACM SIGPLAN Notices , Proceedings of the fifth international conference on Architectural support for programming languages and operating systems ASPLOS-V**, Volume 27 Issue 9

Publisher: ACM Press

Full text available:  pdf(1.33 MB) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



32 Processor coupling: integrating compile time and runtime scheduling for parallelism

 Stephem W. Keckler, William J. Dally

April 1992 **ACM SIGARCH Computer Architecture News , Proceedings of the 19th annual international symposium on Computer architecture ISCA '92**, Volume 20 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.32 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

The technology to implement a single-chip node composed of 4 high-performance floating-point ALUs will be available by 1995. This paper presents processor coupling, a mechanism for controlling multiple ALUs to exploit both instruction-level and inter-thread parallelism, by using compile time and runtime scheduling. The compiler statically schedules individual threads to discover available intra-thread instruction-level parallelism. The runtime scheduling mechanism interleaves threads, explo ...



33 GT-EP: a novel high-performance real-time architecture

 Wei Siong Tan, H. Russ, Cecil O. Alford

April 1991 **ACM SIGARCH Computer Architecture News , Proceedings of the 18th annual international symposium on Computer architecture ISCA '91**, Volume 19 Issue 3

Publisher: ACM Press

Full text available:  pdf(1.09 MB) Additional Information: [full citation](#), [references](#), [index terms](#)



34 A parallel pipelined processor with conditional instruction execution

 Rod Adams, Gordon Steven

March 1991 **ACM SIGARCH Computer Architecture News**, Volume 19 Issue 1

Publisher: ACM Press

Full text available:  pdf(602.87 KB) Additional Information: [full citation](#), [abstract](#), [index terms](#)



In a recent paper by Smith, Lam and Horowitz [1] the concept of 'boosting' was introduced, where instructions from one of the possible instruction streams following a conditional branch were scheduled by the compiler for execution in the basic block containing the branch itself. This paper describes how code from both instruction streams following a conditional branch can be considered for execution in the basic block containing the branch. Branch conditions are stored in Boolean registers and a ...

35 The design of a RISC based multiprocessor chip

Rajiv Gupta, Michael Epstein, Michael Whelan

November 1990 **Proceedings of the 1990 ACM/IEEE conference on Supercomputing Supercomputing '90**

Publisher: IEEE Computer Society

Full text available:  pdf(1.10 MB) Additional Information: [full citation](#), [abstract](#), [references](#)



This paper describes the architecture of a RISC based multiprocessor chip. The processors operate in a MIMD fashion executing parallel instruction streams generated by a parallelizing compiler for the exploitation of fine-grained parallelism. Low cost synchronization mechanisms are supported in hardware. The resulting system is tolerant of unpredictable delays in the progress of individual streams. Instruction level parallelism is exploited through the use of register channels and a mechanism f ...

Keywords: collective branching, fuzzy barrier, parallelizing compiler, register channels, very long instruction word (VLIW) architectures

36 The case for the sustained performance computer architecture

 Apostolos Dollar, Robert F. Krick

December 1989 **ACM SIGARCH Computer Architecture News**, Volume 17 Issue 6

Publisher: ACM Press

Full text available:  pdf(613.67 KB) Additional Information: [full citation](#), [citations](#), [index terms](#)



37 Available instruction-level parallelism for superscalar and superpipelined machines

 N. P. Jouppi, D. W. Wall

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the third international conference on Architectural support for programming languages and operating systems ASPLOS-III**, Volume 17 Issue 2

Publisher: ACM Press

Full text available:  pdf(1.38 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



Superscalar machines can issue several instructions per cycle. Superpipelined machines can issue only one instruction per cycle, but they have cycle times shorter than the latency of any functional unit. In this paper these two techniques are shown to be roughly equivalent ways of exploiting instruction-level parallelism. A parameterizable code reorganization and simulation system was developed and used to measure instruction-level parallelism for a series of benchmarks. Results of these si ...

38 SIMP (Single Instruction stream/Multiple instruction Pipelining): a novel high-speed single-processor architecture

 K. Murakami, N. Irie, S. Tomita

April 1989 **ACM SIGARCH Computer Architecture News , Proceedings of the 16th annual international symposium on Computer architecture ISCA '89**, Volume 17 Issue 3

Publisher: ACM Press

Full text available:  pdf(1.23 MB) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



SIMP is a novel multiple instruction-pipeline parallel architecture. It is targeted for enhancing the performance of SISD processors drastically by exploiting both temporal and spatial parallelisms, and for keeping program compatibility as well. Degree of performance enhancement achieved by SIMP depends on; i) how to supply multiple instructions continuously, and ii) how to resolve data and control dependencies effectively. We have devised the outstanding techniques for instruction fetch an ...

39 The white dwarf: a high-performance application-specific processor

 A. Wolfe, M. Breternitz, C. Stephens, A. L. Ting, D. B. Kirk, R. P. Bianchini, J. P. Shen

May 1988 **ACM SIGARCH Computer Architecture News , Proceedings of the 15th Annual International Symposium on Computer architecture ISCA '88**, Volume 16 Issue 2



Publisher: IEEE Computer Society Press, ACM Press

Full text available:  pdf(1.40 MB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

This paper presents the design and implementation of a high-performance special-purpose processor, called The White Dwarf, for accelerating finite element analysis algorithms. The White Dwarf CPU contains two Am29325 32-bit floating-point processors and one Am29332 32-bit ALU, and employs a wide-instruction word architecture in which the application algorithm is directly implemented in microcode. The entire system is VME-bus compatible and interfaces with a SUN 31160 host. The syste ...

40 Exploiting horizontal and vertical concurrency via the HPSm microprocessor



 Wen-Mei W. Hwu, Yale N. Patt

December 1987 **Proceedings of the 20th annual workshop on Microprogramming MICRO 20**

Publisher: ACM Press

Full text available:  pdf(801.30 KB)

Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#)

HPSm is a single-chip microarchitecture designed and implemented at the University of California to achieve high performance. The approach is to exploit both vertical and horizontal concurrency in the microarchitecture. Experiments have been conducted to demonstrate the effectiveness of HPSm as compared to a popular single-chip microarchitecture, the Berkeley RISC/SPUR. Evaluations have been done with both control intensive and floating point intensive benchmarks. For both types of benchmar ...

Results 21 - 40 of 44

Result page: [previous](#) [1](#) [2](#) [3](#) [next](#)

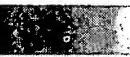
The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)


[Subscribe \(Full Service\)](#) [Register \(Limited Service, Free\)](#) [Login](#)
 The ACM Digital Library The Guide

 +"very long instruction word" +simulat* +parallel +"pipeline stage"

[The ACM DIGITAL LIBRARY](#)

[Feedback](#) [Report a problem](#) [Satisfaction survey](#)

Terms used: **very long instruction word simulat parallel pipeline stage**

Found 44 of 204,472

 Sort results by
 [Save results to a Binder](#)
[Try an Advanced Search](#)

 Display results
 [Search Tips](#)
[Try this search in The ACM Guide](#)
 [Open results in a new window](#)

Results 41 - 44 of 44

 Result page: [previous](#) [1](#) [2](#) [3](#)

Relevance scale

41 A VLIW architecture for a trace scheduling compiler

 Robert P. Colwell, Robert P. Nix, John J. O'Donnell, David B. Papworth, Paul K. Rodman
October 1987 **ACM SIGARCH Computer Architecture News , ACM SIGPLAN Notices , ACM SIGOPS Operating Systems Review , Proceedings of the second international conference on Architectual support for programming languages and operating systems ASPLOS-II**, Volume 15 , 22 , 21 Issue 5 , 10 , 4

Publisher: IEEE Computer Society Press, ACM Press

Full text available:  [pdf\(1.59 MB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)

Very Long Instruction Word (VLIW) architectures were promised to deliver far more than the factor of two or three that current architectures achieve from overlapped execution. Using a new type of compiler which compacts ordinary sequential code into long instruction words, a VLIW machine was expected to provide from ten to thirty times the performance of a more conventional machine built of the same implementation technology. Multiflow Computer, Inc., has now built a VLIW called the TRACE™< ...



42 Evon: an extended von Neumann model for parallel processing

Wai-Mee Ching
November 1986 **Proceedings of 1986 ACM Fall joint computer conference ACM '86**

Publisher: IEEE Computer Society Press

Full text available:  [pdf\(1.01 MB\)](#) Additional Information: [full citation](#), [references](#), [citations](#), [index terms](#)



43 Optimal pipelining in supercomputers

 S. R. Kunkel, J. E. Smith
June 1986 **ACM SIGARCH Computer Architecture News , Proceedings of the 13th annual international symposium on Computer architecture ISCA '86**, Volume 14 Issue 2

Publisher: IEEE Computer Society Press, ACM Press

Full text available:  [pdf\(774.34 KB\)](#) Additional Information: [full citation](#), [abstract](#), [references](#), [citations](#), [index terms](#)



This paper examines the relationship between the degree of central processor pipelining and performance. This relationship is studied in the context of modern supercomputers. Limitations due to instruction dependencies are studied via simulations of the CRAY-1S. Both scalar and vector code are studied. This study shows that instruction dependencies

severely limit performance for scalar code as well as overall performance. The effects of latch overhead are then considered. The prim ...

44 PIPE: a VLSI decoupled architecture

 J. R. Goodman, Jian-tu Hsieh, Koujuch Liou, Andrew R. Pleszkun, P. B. Schechter, Honesty C. Young

June 1985 **ACM SIGARCH Computer Architecture News , Proceedings of the 12th annual international symposium on Computer architecture ISCA '85**, Volume 13 Issue 3

Publisher: IEEE Computer Society Press, ACM Press

Full text available:  [pdf\(770.55 KB\)](#) Additional Information: [full citation](#), [citations](#), [index terms](#)



Results 41 - 44 of 44

Result page: [previous](#) [1](#) [2](#) [3](#)

The ACM Portal is published by the Association for Computing Machinery. Copyright © 2007 ACM, Inc.

[Terms of Usage](#) [Privacy Policy](#) [Code of Ethics](#) [Contact Us](#)

Useful downloads:  [Adobe Acrobat](#)  [QuickTime](#)  [Windows Media Player](#)  [Real Player](#)